arXiv:0910.4135v1 [stat.ML] 21 Oct 2009

Sparsification and feature selection by compressive linear regression

Florin Popescu Ph.D.* and Daniel Renz,
*Fraunhofer FIRST, Intelligent Data Analysis - IDA,
Kekuléstraße 7, 12489 Berlin, Germany

Abstract

The Minimum Description Length (MDL) principle states that the optimal model for a given data set is that which compresses it best. Due to practial limitations the model can be restricted to a class such as linear regression models, which we address in this study. As in other formulations such as the LASSO and forward step-wise regression we are interested in sparsifying the feature set while preserving generalization ability. We derive a well-principled set of codes for both parameters and error residuals along with smooth approximations to lengths of these codes as to allow gradient-descent optimization of description length, and go on to show that sparsification and feature selection using our approach is faster than the LASSO on several datasets from the UCI and StatLib repositories, with favorable generalization accuracy, while being fully automatic, requiring neither cross-validation nor tuning of regularization hyper-parameters, allowing even for a nonlinear expansion of the feature set followed by sparsification.

1 Introduction

1.1 MDL

The Minimum Description Length (MDL) principle is one of many methods that have been proposed to tackle the general model selection problem through principled balancing of the cost of storing a given model's prediction errors on a given data set and that of storing the model's parameters. Based on algorithmic information theory, it was first formulated by Rissanen in 1978 [15]. Other methods include AIC [1], BIC [20] and MML [27], which are based on Bayesian statistics and classical information theory. In the past decades several theoretical approaches to MDL have evolved [12, 15, 26], while practical implementations remain an open research topic [19], as the potential computational burden MDL can impose is significant.

Practical MDL uses description methods that are less expressive than universal languages, so that the length of the shortest description of the data is always computable [13, 19, 21, 22]. It restricts the set of allowed codes in a manner which allows us to find the shortest code length of the data, relative to the set of allowed codes in finite time. So far, three practical MDL schemes have been devised [11]. We used the simplest of them, which is called *two-part MDL*: $\hat{M} = \text{argmin}_{M \in \mathcal{M}} \left(L\left(D|M\right) + L\left(M\right) \right)$, where D is the data, M the model, $\mathcal{M}$ the model class (the set of allowed models) and L the code length function. If we knew the probability distributions $p(M), p(D|M)$, we could use maximum likelihood estimation (the problem being in essence a Bayesian learning formulation) or appropriate asymptotically compact codes, i.e. Shannon-Fano coding, to generate the code books for $L(D|M)$ and $L(M)$ (as it is done in MML). However, we do not know them, so in order to cover a wide

set of possible underlying noise probability densities and size of training data sets we used a novel approach based on codes we developed ourselves.

1.2 Regularized linear regression vs. compressive linear regression

In linear regression, we assume we have data organized in an observation matrix $\mathrm{D} \in \mathbb{R}^{J \times \mathrm{N}}$ where N is the number of observations and J is the number of features. This matrix may be expanded in two ways: by adding functions of groups of observations (therefore increasing N), as in kernel expansion, or adding functions of feature groups (feature expansion), thereby increasing J. In linear regression it is customary to add at least one constant row to the features, corresponding to the bias. Feature expansion can be more generally written as

$$\mathrm{X} \in \left\{ \{\mathbf{b}_q, f_q : q = 1, 2, 3, ..., p\} : \mathbf{b}_q \subset \{1, 2, 3, ..., J\}, f_q : \mathbb{R}^{|\mathbf{b}_q|} \rightarrow \mathbb{R}^1 \right\} \tag{1}$$

Thus X, which we call the *feature product*, is a subset of matrix functions $\mathrm{X} \subset \mathbb{R}^{J \times N} \rightarrow \mathbb{R}^{(J+p) \times N}$ operating on groups of rows, as to allow eventual feature selection. The linear regression formulation can be written as

$$\boldsymbol{y} = X(D)\boldsymbol{\theta} + \boldsymbol{e}, \boldsymbol{y} \in \mathbb{R}^N, \boldsymbol{\theta} \in \mathbb{R}^K, \tag{2}$$

where $K = J+p$. Depending on our prior knowledge of the probability density function $p(\boldsymbol{e})$ (i.e. of measurement errors) there are various ways in which an approximate solution $\hat{\boldsymbol{\theta}}$ may be sought, by minimizing the appropriate loss function (or log-likelihood). One of the most common "non-uniformative" assumptions, consistent with a maximal entropy distribution for finite variance i.i.d sampling, is that the errors are Gaussian and the corresponding loss function to be minimized is provided by the 2-norm

$$\hat{\boldsymbol{\theta}} = \arg\min_{\boldsymbol{\theta}} \left(\|\boldsymbol{y} - X(D)\boldsymbol{\theta}\|_2 + \gamma \|\boldsymbol{\theta}\|_r \right) \tag{3}$$

The additional r-norm is necessary, even without explicit regularization ($\gamma = 1$), to cover the underdetermined case ($K > N$), but without regularization ℓ_2 regression can very often overfit [8]. Various regularization norms have been proposed, the most commonly used being ridge regression and LASSO [23]. The Tikhonov regularization [25] proposes a generalized 2-norm, but commonly the straightforward ℓ_2 norm is used (ridge regression). The LASSO proposes the ℓ_1 norm and offers the additional advantage of a tendency to sparsify $\boldsymbol{\theta}$. Both approaches can be seen, from a Bayesian perspective, to assume a Gaussian error distribution $p(\boldsymbol{e})$ and independent zero-mean priors with uniform variances over components of $\boldsymbol{\theta}$, of Gaussian form for ridge regression and double-exponential for the LASSO. The Bayesian interpretation is problematic however: each feature may have different measurement units (scale), so a likelihood interpretation based on the sum of absolute values of the coefficients (each having inverse units of the corresponding feature) is rather non-informative. Ridge-regression is likewise not scale invariant. Furthermore, the variance of the parameter "prior" depends on the hyper-parameter γ, which is almost always chosen via posterior analysis of the data. Although LASSO and ridge regression, for fixed γ, consist of convex optimization problems which can be accurately and efficiently solved, they are not convex in γ: complete model selection requires computationally expensive bootstrapping and cross-validation, while the actual statistic to be minimized depends on the particular implementation.

Our objective was to provide for a scale invariant objective function which allows for meaningful Bayesian interpretation and does not require hyper-parameter tuning. The MDL formulation of regularized regression requires, first of all, the weak assumption that the target

has finite and uniform measurement precision or quantization width $\delta(\boldsymbol{y})$ (i.e. $\boldsymbol{y}/\delta(\boldsymbol{y}) \in \mathbb{Z}^N$). We seek

$$\min_{\boldsymbol{\theta}_\#} \left(L\left(\frac{\boldsymbol{y}}{\delta(\boldsymbol{y})} - \left\lfloor \frac{X(D)\boldsymbol{\theta}_\#}{\delta(\boldsymbol{y})} \right\rfloor \right) + L\left(\boldsymbol{\theta}_\#\right) \right), \text{ where } \boldsymbol{\theta}_\# \in \mathbb{Q}^K \tag{4}$$

The difference between equations (3) and (4) is that we are optimizing over rational-valued vectors $\boldsymbol{\theta}_\#$, rather than real-valued ones, as both the parameter prior and the measurement error distributions are now discrete probabilities rather than probability density functions. Clearly, a direct optimization over rational numbers would be very difficult. This is precisely why we develop codes, approximations and bounds to make MDL optimization possible for the model class of linear regression models with nonlinear feature products. We shall call this method *Compressive Linear Regression (CLR)*.

2 Methods

The basic approach we will follow is to compute, as precisely as possible, description lengths of regression parameters and residuals such that a smooth objective function which approximates the description length of the data can be built. For this purpose, known problems such as compact codes for integers, rationals and random integer sequences must be revisited.

2.1 Coding of integers

Rissanen's universal prefix-free code (that is, no code word is a prefix of any other code word) for integers [17] provides a Bayesian "non-uniformative" prior $2^{-L(i)}$, a principle which can be extended from the integer set to rational numbers as well, which we will propose in this paper. A universal code is a prefix-free code which has an expected code length over all integers that is invariant, within a constant, no matter what the prior probability of the integers are (given that the latter monotonically decreases). However, Rissanen's code is not very compact as defined by the Kraft inequality: $\sum_{d\in\mathcal{D}} 2^{-L(d)} \leq 1$. A *compact code* is a code, for which the Kraft sum is close to one, which implies that only few of all possible code words of a certain length are not in the code. For every set of code word lengths that satisfy the inequality, there exists a unique decodable prefix-free code with the same code word lengths [9]. This is why we can restrict our attention to prefix-free codes throughout this paper. Another argument against using Rissanen's code is that the area under the curve of the code length function is not convex everywhere, being concave at $x_1 = 2, x_2 = 2^2 = 4, x_3 = 2^4 = 16, x_4 = 2^{16}, ...$, which can cause problems for some optimization algorithms.

Universal codes are the basic building block for the other codes described, therefore we chose to build a smooth, compact and convex universal code U_n for non-negative integers (and consequently for the signed integers as well via the ordering 0,-1,1,-2,2,-3 etc. We call the code for the signed integers U). Our code uses two different coding schemes: For small numbers, it uses a coding scheme F similar to Fibonacci coding [2] and then switches at a certain point to a coding scheme E producing code words with the same lengths as Elias-Delta coding [7] for large numbers. In F the first two codewords are 0 and 100. The following k codewords are produced by binary-adding the number 1. For k we use consecutive Fibonacci numbers, i.e. 1,1,2,3,5,... The k-th codeword is additionally appended by "0". In E, the codeword for a number n consists of the three binaries t, $blen$ and b^*. t are $len - 1$ ones, len is the length of the binary representation $b(n)$ of n, $blen$ is the binary representation of $b(n)$ with the first bit set to zero and b^* is $b(n)$ without the first bit.

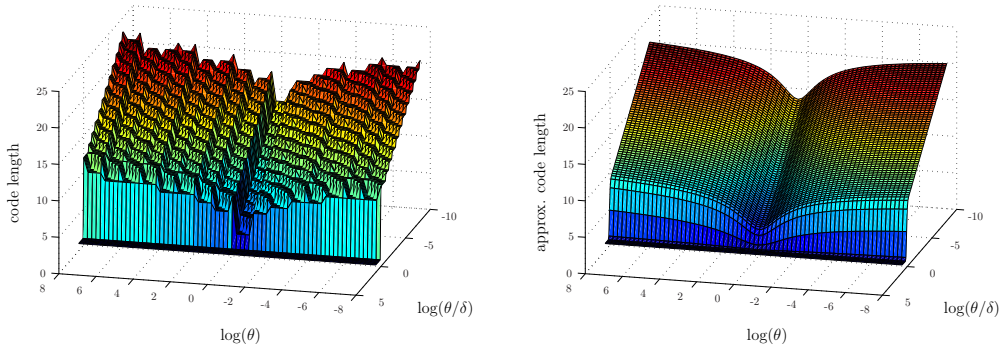

Figure 1: Left: α code for rational numbers: Code word lengths increase with $|\theta|$ and θ/δ. Right: Smooth approximation function for the α coding function.

2.2 Coding of rationals or simplest rational within an interval on the real line

Any universal integer code can be used to construct a prefix-free code for rational numbers of specified precision. However, there is no obvious ordering scheme for rationals in terms of code lengths. Li and Vitányi offer a method to map rationals into the unit interval $[0, 1] \in \mathbb{R}$ using the idea of cylinder sets [13]. We implement a variant of this code and name it α-code:

$$\alpha : \mathbb{R}^2 \to \{0,1\}^* \text{ , } \alpha^{-1}\text{: } \{0,1\}^* \to \mathbb{Q} \text{ , s.t. } \theta_\# = \alpha^{-1}\left(\alpha(\theta,\delta)\right) \in (\theta-\delta, \theta-\delta)) \tag{5}$$

where $\alpha^{-1}(\alpha(\theta,\delta))$ is the rational number coded by the binary $\alpha(\theta,\delta)$. Using the code α, we can recover a rational number close to the *real* number θ to a precision of at least δ. We define $\alpha = \alpha(\theta,\delta) = U(\theta_\delta)U(\theta_{\log})$, where θ is the rational number to be encoded, δ the precision to which it should be encoded, $\theta_\delta = \lceil \theta/\delta \rceil$ and $\theta_{\log} = \lceil \log|\theta| \rceil - \lceil 0.5 \log n \rceil$. The decimal value of the code word $\alpha(\theta,\delta)$ is $\theta_\# = \theta_\delta \cdot 2^{\lceil \log \theta_\delta \rceil + \theta_{\log}}$. As α is prefix-free, the code for a vector of rationals which are not presumed to be interdependent is simply the concatenation of the codes of its elements: $\alpha(\boldsymbol{\theta}, \boldsymbol{\delta}) = \sum_i \alpha(\theta_i, \delta_i)$. We shall use α to store the linear regression parameters, as detailed below. Through numerical tests, we were able to find a smooth function $\bar{\alpha}$ that approximates our α-code lengths, for which we calculated an approximation mean error of 0.8 bit by sampling 10^5 evenly spaced points for $\theta \in [2^{-8}; 2^8]$ and a *relative precision* $\theta/\delta \in [2^{-8}; 2^0]$:

$$\begin{aligned} \bar{\alpha} &= c_0 \cdot \log\left(\tau(\theta) \cdot \left(\operatorname{erf}\left(10 \cdot \left(\frac{\theta-\delta}{\theta}\right) + 1\right) \cdot |\theta| + \delta\right) - \log\delta\right) + 1 \\ \tau(\theta) &= \left|c_1 \cdot \left(\log\left(\log\left(\theta^2 + c_2\right)\right)\right) - 1\right| \end{aligned} \tag{6}$$

In this formula, `erf` is the error function and c_0, c_1 and c_2 are numerically fitted constants of O(1). The alpha code lengths and their smooth approximations are shown in Figure 1.

2.3 Spherical coding of a random sequence of signed integers

We seek a means of coding the part of the data that cannot be explained by the model, i.e. the residuals $\boldsymbol{e}$ (a sequence of signed integers). Since only model fitting (compression) can determine that the sequence is not random, we shall build a code assuming that each element is random (incompressible), independent of each other, and uniform (equivalent to an i.i.d sample).

One way we could achieve this is simply to assign a uniform length signed integer code to all possible values ranging from the observed maximum to the observed minimum. This would be equivalent to assuming a uniform error distribution, which would be inefficient as it is not likely in practice. Another possibility would be to assign a universal code to each signed integer, which would be also inefficient, since we would not be exploiting the (reasonable) assumption of uniformity. We chose, therefore, to assume that the likelihood of a sample $\boldsymbol{e}$ decreases with its 2-norm (hence the name "spherical code"), and to use a universal code for coding of each possible N-dimensional value in order of increasing 2-norm. This, we will show numerically at least, is equivalent to assuming a Gaussian (normal) distribution of noise.

The coding problem, i.e. a scheme which allows the counting of errors of any size, becomes one of counting the maximal expected number of hypercubes with side lengths $\delta(X)$ that intersect a hypersphere of N dimensions ordered spirally outwards from the origin.

$$H : N \leftrightarrow \mathbb{Z}^N, \;\; s.t. \left\| H(n) \right\|_2 \leq \left\| H(n+1) \right\|_2 \;\; \forall n \in N \tag{7}$$

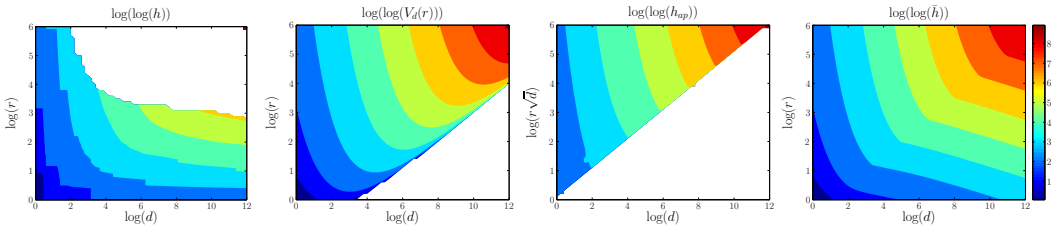

Figure 2: Spherical counting scheme and approximations: The first picture shows the surface plot of the exact count h. White indicates the region in which computation of h was not feasible. In the second picture we can see the spherical volume. In the lower right region this approximation is not accurate. The third picture shows the Shannon coding length for a spherical data distribution, where we plot d vs. $r/\sqrt{d}$. The reason why this approximation is not valid in the region underneath the first angle bisecting line is that $r/\sqrt{d}$ gets so small that the coding length would be estimated as one bit. The fourth picture shows the approximation function $\bar{h}$.

So, our spiral counting scheme $H(n)$ guarantees that the 2-norm is monotonically increasing and therefore, since the universal description length of an integer is monotonically increasing with the integer, it follows that under this counting scheme the description length of an integer vector under H is monotonically increasing with its 2-norm. This leads us to state that:

$$\left\| \frac{\boldsymbol{y}}{\delta(X)} \right\|_2 \leq r \Rightarrow K(\boldsymbol{y}) = \left| U\left(H^{-1}\left(\frac{\boldsymbol{y}}{\delta(X)} \right) \right) \right| \leq h(r,N) \lessapprox \left| U\left(\max\left(V_S^N(r), 1 \right) \right) \right|, \tag{8}$$

where $V_S^N(r)$ denotes the volume of the hyper-sphere of radius r in $\mathbb{R}^N$. We can establish a connection between h and the Shannon entropy $h_{sh}(x)$. It is the shortest average coding length for a random variable x of dimension N. For spherical data with variance σ^2,

$$h_{sh}(x) \;\; = \;\; \frac{N}{2} \cdot \log(2\pi e) + N \cdot \log(\sigma) \tag{9}$$

To be able to compare h_{sh} and h, we have to consider the case of finite data again: We cannot simply assume that σ is known, but rather we have to actually encode it and consider

its coding length. We do so using the α-code. The minimum number of bits that we need to store σ can be found by minimizing over the precision $\Delta\sigma$. We call the resulting *applied* Shannon coding length h_{ap}:

$$h_{ap}(x) = \frac{N}{2} \cdot \log(2\pi e) + \min_{\Delta\sigma}(N \cdot \log(\sigma + \Delta\sigma) + \alpha(\sigma, \Delta\sigma)) \approx h(r = \sigma\sqrt{N}, N) \tag{10}$$

Figure 2 shows that the applied Shannon coding length is approximated by the spherical code length function for large N. Shannon-Fano or Huffman code lengths/entropies for a given distribution are only valid asymptotically, i.e. for very large N. For finite, small N, such coding - based on known or sample standard deviation (SD) - may even result in expansion rather than the intended compression.

2.4 MDL for linear regression

The complete (loss-less) code length approximation λ_M of the data to be minimized is:

$$\lambda_M(\boldsymbol{\theta}, \boldsymbol{\delta}, X) = \bar{\alpha}(\boldsymbol{\theta}, \boldsymbol{\delta}) + \bar{h}(s^2_M(\boldsymbol{\theta}_{\#}, X), N), \tag{11}$$

where $\boldsymbol{\theta}_{\#} = \alpha^{-1}(\alpha(\boldsymbol{\theta}, \boldsymbol{\delta}))$. Bars on α and h denote smooth approximations. Note that this equation requires the function $s_M(\cdot)$, which relates the size of the residual 2-norm to the parameters and their storage precision, for some model M which specifies how the parameters encode $\boldsymbol{y}$, usually represented by a Hessian function. This principle can be applied to any type of regression fit: we shall do so for linear regression.

Finding a close, smooth bound $\bar{h}$ for any r and N is a difficult mathematical problem, which also can be formulated as finding the number of representations of an integer as a sum of squares. A general (approximative) function has not been found yet [10]. For combinations of large enough r (> 5) and low enough N ($< 10^6$) (see previous section) the following approximation based on the unit sphere volume is reasonable:

$$\bar{h}(s^2(\boldsymbol{e})) = \log\left(\frac{\pi^{\frac{N}{2}}}{\Gamma\left(\frac{N}{2}+1\right)}\right) + \frac{N}{2}\log\left(\frac{s^2(\boldsymbol{\theta}_{\#}, X)}{\delta(X)^2}\right) \tag{12}$$

The α code has a rather nice property in that the numerical value of the decoding $\alpha^{-1}(\alpha(\theta_i, \delta_i)) = \theta_i + \delta(\theta_i)$ is uniformly distributed on the interval $[\theta_i - \delta_i, \theta_i + \delta_i]$ with $p_i(\delta) = (2\delta_i)^{-1}$, over all possible pairs of θ and δ: as Monte-Carlo simulations have confirmed, the code behaves without bias of $((\theta_{\#})_i - \theta_i)/\delta_i$. This allows us to take the expected value of the increase in the norm of the residual for random, independent perturbations $\delta(\theta_i)$ of each θ_i with the uniform probability $p_i(\delta)$. Calculating the growth of residual norm (squared) with respect to parameter accuracy $\boldsymbol{\delta}$ is straightforward for linear regression:

$$\begin{aligned} s^2(\alpha^{-1}(\alpha(\boldsymbol{\theta}, \boldsymbol{\delta})), X) &= \left(\boldsymbol{y} - X \cdot \alpha^{-1}(\alpha(\boldsymbol{\theta}, \boldsymbol{\delta}))\right)^T \cdot \left(\boldsymbol{y} - X \cdot \alpha^{-1}(\alpha(\boldsymbol{\theta}, \boldsymbol{\delta}))\right) \\ &= \left(\boldsymbol{y} - X \cdot (\boldsymbol{\theta} + \boldsymbol{\delta}(\boldsymbol{\theta}))\right)^T \cdot \left(\boldsymbol{y} - X \cdot (\boldsymbol{\theta} + \boldsymbol{\delta}(\boldsymbol{\theta}))\right) \end{aligned} \tag{13}$$

We choose $\boldsymbol{\theta} = \left(X^T X\right)^{-1} X^T y$ to minimize the 2-norm of the residual $\boldsymbol{e}^T\boldsymbol{e}$, (or any exact solution in the underdetermined case) so the gradient is with respect to $\boldsymbol{\theta}$ zero, meaning

that any perturbation from $\boldsymbol{\theta}$ results in strictly quadratic growth: $\boldsymbol{e}^T\boldsymbol{e}+\boldsymbol{\delta}(\boldsymbol{\theta})^T X^T X\boldsymbol{\delta}(\boldsymbol{\theta}) = \boldsymbol{e}^T\boldsymbol{e} + \boldsymbol{\delta}(\boldsymbol{\theta})^T\Sigma_X\boldsymbol{\delta}(\boldsymbol{\theta})$, where Σ_X is the covariance of X. Taking the expected value over possible perturbations (i.e. the outputs of the coder-decoder):

$$\begin{aligned}
E\left[\delta(\theta_i)^T(\Sigma_X)_{i,i}\delta(\theta_i)\right] &= \int_{-\delta_i}^{\delta_i} p(\delta(\theta_i))(\Sigma_X)_{i,i}\delta(\theta_i)^2 d\delta(\theta_i) \\
&= (2\delta_i)^{-1}\frac{2}{3}(\Sigma_X)_{i,i}\delta_i^3 = \frac{1}{3}(\Sigma_X)_{i,i}\delta_i^2 \\
E\left[\delta(\theta_i)^T(\Sigma_X)_{i,j}\delta(\theta_j)\right]_{i\neq j} &= c\int_{-\delta_j}^{\delta_j}\int_{-\delta_i}^{\delta_i}(\Sigma_X)_{i,j}\delta(\theta_i)\delta(\theta_j)d\delta(\theta_i)d\delta(\theta_j) = 0 \\
E\left[\delta(\boldsymbol{\theta})^T\Sigma_X\delta(\boldsymbol{\theta})\right] &= E\left[\sum_{i,j}\delta(\theta_i)^T(\Sigma_X)_{i,j}\delta(\theta_j)\right] = \frac{1}{3}\sum_i(\Sigma_X)_{i,i}\delta_i^2 \qquad (14)
\end{aligned}$$

Finally, a differentiable CLR objective function for linear regression can be written.

$$\begin{aligned}
\min_{\boldsymbol{\theta}}\left(E_{\delta(\theta)}\left[\lambda_M(\boldsymbol{\theta},\boldsymbol{\delta},X)\right]\right) &= \min_{\boldsymbol{\theta}}\left(\bar{\alpha}(\boldsymbol{\theta},\boldsymbol{\delta}) + E_{\delta(\theta)}\left[\bar{h}(s_M^2+\delta(X)^T\Sigma_{X,M}\delta(X),N)\right]\right) \\
&\approx \min_{\boldsymbol{\theta}}\left(\sum_i\bar{\alpha}(\theta_i,\delta_i) + \bar{h}(s_M^2+\frac{1}{3}\sum_i(\Sigma_X)_{i,i}\delta_i^2,N)\right) \qquad (15)
\end{aligned}$$

The approximation $E[h(a+x)] = E[h(a)+h'(a)x+...] = h(a)+h'(a)E[x]+... \cong h(a+E[x])$ is valid if $h(a)$ is locally linear - being $O(\log(a))$ that assumption is reasonable. The resulting objective function is smooth in optimization parameters ($2K$ such parameters: $\boldsymbol{\theta}$ and $\boldsymbol{\delta}$) but is non-convex. A useful unbiased heuristic we have found was to use simplex optimization (code was written in Matlab, Mathworks, Inc.) with starting value for $\boldsymbol{\theta}$ at the usual 2-norm solution and $\delta_i = |\theta_i|/2$. After each downhill optimization, all parameters for which $\delta_i > |\theta_i|$ were discarded and the entire procedure repeated (with reduced X) until no further parameters were thus "'culled"'.

3 Results

The CLR method proposed was compared to 2 different LASSO variants on both simulated and real data. Simulated data allows us to gauge the relative ability of CLR to accurately recover the sparsity structure in the case in which it is known. We produced the same synthetic data sets as did Tibshirani in his seminal LASSO paper [23]. For each of three examples, we simulated 50 data sets consisting of 20 observations from the model $y = \boldsymbol{\theta}^T\boldsymbol{x}+\sigma\boldsymbol{\epsilon}$, where ϵ is standard normal. The correlation between x_i and x_j was $\rho^{|i-j|}$ with $\rho = 0.5$. The dataset paramters for SIM1 were $\boldsymbol{\beta} = (3, 1.5, 0, 0, 2, 0, 0, 0)^T$ and $\sigma = 3$, for SIM2 $\beta_j = 0.85, \forall j$ and $\sigma = 3$ and for SIM3 $\beta = (5, 0, 0, 0, 0, 0, 0, 0)$ and $\sigma = 2$. In [23] LASSO performed favorably compared to other methods such as ridge regression and therefore only LASSO and normal 2-norm regression was performed for comparison. Performance results are given in Table 1. for each of the simulated datasets (SIM1, SIM2, SIM3) and methods evaluated (CLR - parsimonious linear regression, LASSOLCR - LASSO largest consistent region, LASSOCV - LASSO with cross-validation, similar to Tibshirani's cross-validated LASSO in [23] which was found to be highly performant). Both variants used bootstrapping for LASSO statistic estimation. We used the LASSO implementation available in the Matlab repository (www.mathworks.com/matlabcentral/fileexchange) written by M. Dunham and described in [4]. The given values are mean running time (as a fraction of L2 running time)

over all simulations, and for each dataset: the number of non-zero parameters computed, the minimum squared error (MSE) of the computed parameters vs. simulated parameters and the square root of the MSE of residuals as a fraction of the size of the noise used for the simulation σ. The bias value is not included in these calculations. For L2 the full parameters set is always returned and therefore 8 vs. 2 means that 2 of the actual parameters for that dataset were non-zero.

Table 1: Results of sparse linear regressions on simulated datasets.

Method		**CLR**	**LASSOLCR**	**LASSOCV**	**L2**
Time		112	204	202	1
nr. of parameters found	**SIM1**	2.10±1.4	0.32±0.5	3.40±2.2	8 vs. 3
	SIM2	1.26±1.1	0.04±0.2	4.08±1.7	8 vs. 8
	SIM3	1.72±1.3	0.90±0.3	2.68±2.0	8 vs. 1
MSE	**SIM1**	1.27±0.8	1.72±0.4	1.01±0.6	1.10±0.8
	SIM2	1.12±0.4	0.74±0.1	1.10±0.6	1.04±0.8
	SIM3	0.19±0.4	0.33±0.9	0.25±0.4	0.72±0.8
$s_d(\boldsymbol{e})/\sigma$	**SIM1**	1.05±0.3	1.68±0.4	0.91±0.2	0.73±0.1
	SIM2	1.28±0.4	1.77±0.3	0.91±0.2	0.75±0.1
	SIM3	0.94±0.2	1.11±0.4	0.87±0.2	0.75±0.2

In order to test generalization performance, we used the following datasets from the UCI repository [3]: Automobile, Concrete Compressive Strength [29], Forest Fires [5], Housing, Qualitative Structure Activity Relationships (Pyrimidines and Triazines). We also used one more dataset from the Statlib Archives (http://lib.stat.cmu.edu/): Detroit [14]. For generalization estimation, we used a training test split of 2/3 to 1/3. A feature product of $\mathrm{X}_{k+1+i} = D_i^2$ was used along with a bias term for all methods. Figure 3 shows training vs. test set errors on the 7 datasets evaluated. The axes are standard deviations of residuals as a fraction of the standard deviation of the target, for both test (x-axis) and training subsets (y-axis). Note that many analyses (especially LASSOLCR) resulted in no features (except bias) being selected, i.e. the ratio of s.d.'s equaled 1. The callout in the top right corner is a list of such instances. The actual test s.d. ratio value for LASSOCV, dataset 4, was 59.0 (off scale). The solution sparsity, averaged over all datasets, was 16.7% for CLR, 11.7% for LASSOLCR and 37.4% for LASSOCV.

4 Conclusion

In this paper, we have outlined both the theoretical justification for a detailed MDL application to linear regression which concurrently provides both regularization, feature selection and model selection. At first glance, the complex derivation may appear to be a bit of an overkill. However, simpler MDL applications to the feature selection problem, which have been previously proposed (MDL as a model-selection is long in use for time-series linear auto-regressive modeling), do not perform as well to our experience. Why so? First of all because code length estimates based on Shannon entropy do not work well for small numbers of examples (hence the meticulous derivation of the spherical code). The dataset results we present herein are precisely for small N problems, where there is not much to compress. In the case of auto-regressive time series modeling, it was previously shown that a simpler version of our method exhibits increasing sparsity structure recovery accuracy with increasing N. The same non-standard MDL/MML situation applies for large values of quantization width δ with respect to standard deviation of the target. In fact, δ can be (and was) inferred directly from the histogram of the target. Histograms of the modelling errors

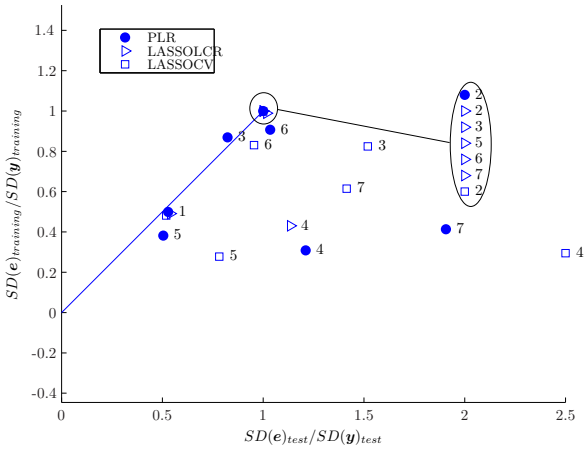

Figure 3: Generalization performance on real data. The datasets are numbered as follows: 1. Concrete (N=1030, J=8), 2. Forest Fires (N=517, J=12), 3. Detroit (N=60, J=15), 4. Housing (N=506, J=13), 5. Automobile (N=159, J=15), 6. Triazines (N=186, J=31), 7. Pyrimidines (N=74, J=26).

(residuals) on the "real" datasets showed that only one of these passed the Jarque-Berra test for normality (Detroit) yet CLR worked well for 2 other datasets. CLR did not perform well on 2 datasets, but neither did either LASSO variant. Note also that LASSO-CV highly overfit on one of the datasets, highlighting the overfit risk of cross-validation. Overall CLR was quite conservative, and it outperformed LASSO on the sparse simulated datasets (in avoiding overfit and recovering the structure). If there is any generally discernible trend in CLR, it is toward underfit, which is not suprising since we are restricting ourselves to a smaller class of models than that from which the 'generating' model may have belonged to.

The CLR method is so conservative, in fact, that it also "works" in the underdetermined case, not only allowing us to use product features such as the square but also all combinations of feature product pairs, allowing for a full multi-dimensional polynomial fit. It may be surprising to note that CLR actually worked faster than the "convex" LASSO although it is not convex. This is because of the reliance on cross-validation and bootstrapping that hyperparameter choice traditionally requires, both in LASSO and other classification/regression paradigms. As the number of features increases, the relative running time of gradient descent algorithms will increase relative to the LASSO, and further work will be necessary to provide for convex approximation to the MDL and customized, more efficient optimization algorithms. As it is known that feature selection is, per se, a non-convex, NP-complete problem [6] it is not likely that the computational time burden will decrease without additional compromises. We'd like to assert, credibly, that there are no current ML approaches which are fully convex, either because of the hyperparameter choice or the adjustment of prior variances. However, our work points to the possibility of fully automatic learning, with no practitioner imposed choices necessary. We have already implemented heuristic algorithms which iteratively generate, and then cull/prune features to minimize MDL, which is an unbounded-time formulation. Since learning in general, at least under the algorithmic information theory perspective, is incomputable, it should not be surprising that heuristic approaches with practically imposed time bounds will be necessary. While the common Gaussian i.i.d assumption was used herein, the MDL formulation allows for learning under other, or even unkown, noise sources. The underlying philosophy of our ML approach is to spend computational resources on model generation and selection rather than validation.

References

[1] H. Akaike. Information theory and an extension of the maximum likelihood principle. *Second International Symposium on Information Theory*, pages 267–281, 1973.

[2] A. Apostolico and A. S. Fraenkel. Robust transmisson of unbounded strings using fibonacci representations. *IEEE Transactions on Information Theory*, IT-33(2):238–245, 1987.

[3] A. Asuncion and D.J. Newman. UCI machine learning repository, 2007.

[4] F. R. Bach. Bolasso: model consistent lasso estimation through the bootstrap. *Proceedings of the 25th international conference on Machine learning*, 307:33–40, 2008.

[5] P. Cortez and A. Morais. A data mining approach to predict forest fires using meteorological data. *J. Neves, M. F. Santos and J. Machado Eds., New Trends in Artificial Intelligence, Proceedings of the 13th EPIA 2007 - Portuguese Conference on Artificial Intelligence*, 2007.

[6] S. Davies and S. Russell. Np-completeness of searches for smallest possible feature sets. In *AAAI Symposium on Intelligent Relevance*, pages 37–39. AAAI Press, 1994.

[7] P. Elias. Universal codeword sets and representations of the integers. *IEEE Transactions on Information Theory*, IT-21(2):194–203, 1975.

[8] G. Golub and C. Van Loan. *Matrix Computation*. John Hopkins University Press, Baltimore, 1989.

[9] P. Grünwald. Model selection based on minimum description length. *Journal of Mathematical Psychology*, 44(1):133–152, 2000.

[10] E. Grosswald. *Representation of integers as sums of squares*. Springer, 1985.

[11] A. D. Lanterman. Schwarz, wallace and rissanen: Intertwining themes in theories of model selection. *International Statistical Review*, 69(2):185–212, 2001.

[12] M. Li and P. M. B. Vitanyi. *Inductive reasoning and Kolmogorov complexity*. Springer-Verlag, 1997.

[13] M. Li and P. M. B. Vitanyi. *An introduction to Kolmogorov complexity and its applications, 2nd edition*. Springer-Verlag, 1997.

[14] R. L. Mason R. F. Gunst. *Regression analysis and its application :a data-oriented approach*. M. Decker, New York, 1980.

[15] J. Rissanen. Modeling by shortest data description. *Automatica*, 14:465–471, 1978.

[16] J. Rissanen. A universal prior for integers and estimation by minimum description length. *Annals of Statistics*, 11(2):416–431, 1983.

[17] J. Rissanen. Universal coding, information, prediction, and estimation. *IEEE Transactions on Information Theory*, IT-30(4):629–636, 1984.

[18] J. Schmidhuber. Discovering neural nets with low kolmogorov complexity and high generalization capability. *Neural Networks*, 10(5):857–873, 1997.

[19] J. Schmidhuber. The speed prior: A new simplicity measure yielding near-optimal computable predictions. In *Proceedings of the 15th Annual Conference on Computational Learning Theory (COLT 2002), Lecture Notes in Artificial Intelligence*, pages 216–228. Springer, 2002.

[20] G. Schwarz. Estimating the dimension of a model. *Annals of Statistics*, 6(2):461–464, 1978.

[21] R. Solomonoff. A formal theory of inductive inference, part i. *Information and Control, Part I*, 7(1):1–22, 1964.

[22] R. Solomonoff. A formal theory of inductive inference, part ii. *Information and Control, Part II*, 7(2):224–254, 1964.

[23] R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, 58(1):267–288, 1996.

[24] R. Tibshirani. The lasso method for variable selection in the cox model. *Statistics in Medicine*, 16:385–395, 1997.

[25] A. Tikhonov and V. Arsenin. *Solution for Ill-Posed Problems*. Wiley, New York, 1977.

[26] P. Vitanyi and M. Li. Minimum description length induction, bayesianism, and kolmogorov complexity. *IEEE Transactions on Information Theory*, 46(2):446–464, 2000.

[27] C. Wallace and P. Freeman. Estimation and inference by compact coding. *Journal of the Royal Statistical Society B*, 49(3):241–252, 1987.

[28] H. Xu, C. Caramanis, and S. Mannor. Robust regression and lasso. 2008.

[29] I. Yeh. Modeling of strength of high performance concrete using artificial neural networks. *Cement and Concrete Research*, 28(12):1797–1808, 1998.

[30] H. Zou and T. Hastie. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society B*, 67:301–320, 2005.

5 Acknowledgements

We thank Prof. Ulf Leser (Humboldt University Berlin) for his helpful comments.

This manuscript was submitted to the Neural Information Processing Conference (NIPS) 2009.